

U.S. Coast Guard Research and Development Center
1082 Shennecossett Road, Groton, CT 06340-6096

Report No. CG-D-03-02

CUTTER CONNECTIVITY BANDWIDTH STUDY



**FINAL REPORT
OCTOBER 2002**



This document is available to the U.S. public through the
National Technical Information Service, Springfield, VA 22161

Prepared for:

U.S. Department of Transportation
United States Coast Guard
System (G-S)
Washington, DC 20593-0001

NOTICE

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

The United States Government does not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to the object of this report.

This report does not constitute a standard, specification, or regulation.



Marc B. Mandler, Ph.D.
Technical Director
United States Coast Guard
Research & Development Center
1082 Shennecossett Road
Groton, CT 06340-6096

Technical Report Documentation Page

1. Report No. CG-D-03-02		2. Government Accession Number		3. Recipient's Catalog No.	
4. Title and Subtitle Cutter Connectivity Bandwidth Study				5. Report Date OCTOBER 2002	
				6. Performing Organization Code Project No. 9270	
7. Author(s)				8. Performing Organization Report No. R&DC 617	
9. Performing Organization Name and Address U.S. Coast Guard Research and Development Center 1082 Shennecossett Road Groton, CT 06340-6096				10. Work Unit No. (TRAIS)	
				11. Contract or Grant No.	
12. Sponsoring Organization Name and Address U.S. Department of Transportation United States Coast Guard Systems (G-S) Washington, DC 20593-0001				13. Type of Report & Period Covered Final Report	
				14. Sponsoring Agency Code Commandant (G-SCT) U.S. Coast Guard Headquarters Washington, DC 20593-0001	
15. Supplementary Notes The R&D Center's technical point of contact is LCDR Brian Krenzien, 860-441-2846, email: bkrenzien@rdc.uscg.mil.					
16. Abstract The goal of this study was to determine how much bandwidth is required for cutters to meet emerging data transfer requirements. The Cutter Connectivity Business Solutions Team with guidance from the Commandant's Innovation Council sponsored this study. Today, many Coast Guard administrative and business functions are being conducted via electronic means. Although our larger cutters can establish part-time connectivity using commercial satellite communications (SATCOM) while underway, there are numerous complaints regarding poor application performance. Additionally, smaller cutters do not have any standard means of underway connectivity. The R&D study shows the most important factor affecting web performance and enterprise applications onboard cutters was latency. Latency describes the time it takes the signal to reach the satellite and come back down through space. The latency due to use of higher orbit satellites is causing poor application performance and inefficient use of expensive SATCOM links. To improve performance, the CG must, (1) reduce latency by using alternate communications links such as low-earth orbit satellites, (2) tailor applications to the SATCOM link and/or (3) optimize protocols used for data communication to minimize time required by present applications to establish communications between the user and the host systems.					
17. Key Words Aggregate Bandwidth, Cutter Connectivity, latency, Satellite latency, OPNET modeling			18. Distribution Statement This document is available to the U.S. public through the National Technical Information Service, Springfield, VA 22161		
19. Security Class (This Report) UNCLASSIFIED		20. Security Class (This Page) UNCLASSIFIED		21. No of Pages	
				22. Price	

Form DOT F 1700.7 (8/72) Reproduction of form and completed page is authorized.

Executive Summary

Currently, most of our capital cutters obtain CGDN+ connectivity through systems such as Inmarsat B, which uses satellites in geosynchronous orbit (35,786 km above the earth's surface) to relay communications signals between ship and shore.

With today's technologies, more of our daily administrative and business functions are being conducted via electronic means. For example, a connection to the intranet (CGDN+) is now required to submit a travel claim using the Unit Travel System (UTS), fill out an e-resume using the Coast Guard Human Resource Management System (CGHRMS), or to access numerous other web-based Enterprise Applications (EAs) including the Coast Guard Message System (CGMS), Large Unit Financial System (LUFS), Abstract of Operations (AOPS) and Marine Information for Safety and Law Enforcement (MISLE). Although most of our larger cutters are now able to establish at least part-time underway CGDN+ connectivity using commercial satellite communications (SATCOM), shipboard personnel are complaining of poor application performance. Importantly, our smaller cutters do not have any standard means of underway connectivity to CGDN+. The goal of this R&D study was to determine how much bandwidth is required for each cutter class to meet e-CG data transfer requirements.

The R&D study shows that increasing the SATCOM bandwidth already available to most large cutters will not necessarily improve web performance or the speed of CG enterprise applications. Network measurements and modeling results showed that the most important factor affecting web performance and enterprise application performance onboard these cutters were latency (the time it takes the signal to propagate through space to reach the satellite and back down). The latency due to the satellite is causing poor application performance and inefficient use of expensive SATCOM links. To improve performance, the CG must, (1) reduce latency by using alternate communications links, (2) tailor our applications for the type of SATCOM link and/or (3) optimize the protocols used for data communication.

Problem Statement

Operational requirements dictate that underway cutters access various CG Enterprise Applications (EAs); however, those requirements have not been translated to bandwidth requirements. Therefore, it is unknown how much commercial SATCOM capacity needs to be leased in order to support the applications onboard various platforms.

Project Objective

The goal of this study was to determine the “aggregate bandwidth” requirements of underway cutters based on throughput measurements of EAs identified as required for underway use. The results of the study will assist CG long-range communications planners in identifying performance gaps and developing an underway communications architecture to resolve them.

Project Background

The USCG Innovation Council chartered a Cutter Connectivity Business Solutions Team (C2BST) in 2001 made up of representatives from G-SCT, G-CIT, G-OCC, TISCOM, LANTAREA, PACAREA and the R&D Center.

“The C2BST was chartered to look at the current state of cutter data connectivity, research & evaluate current and future technologies, and recommend a way ahead for the CG to achieve the Commandant’s e-CG vision. Cutter connectivity is viewed as one of the more difficult obstacles to achieving e-CG. With one of the goals of e-CG being ‘to allow all members the ability to readily go online, anytime, anywhere to not only meet mission requirements and conduct necessary business, but to address personal logistics needs as well (travel, pay, household moves, assignment preference, medical appointments, emergency data, etc.)’ the CG must find a way to get cutters near full time connectivity to the CG Data Network (CGDN+).”

-excerpt from C2BST Report (Aug 2001)

The C2BST reported on current state capabilities to serve as a baseline and identified ongoing cutter connectivity initiatives in order to avoid duplication of effort. The team recommended an Aggregate Bandwidth Study be conducted to provide a standard against which to measure the CG’s existing and future cutter connectivity shortfalls. In addition, several technologies were investigated in order to avert the projected shortfalls.

Information on which EAs would be required underway was derived from interviews with Headquarters Program Managers and Cutter representatives. The cutters were divided into two groups for the purposes of functional requirements:

- (1) Typically underway for *more* than a week at a time
- (2) Typically underway for *less* than a week at a time

Project Description

A product called OPNET IT Guru, manufactured by OPNET Technologies, was used to model the wireless link between ship and shore. OPNET's Application Characterization Environment (ACE) was used to capture traffic for each of the EAs from the Coast Guard Data Network Plus (CGDN+). The data collected from the terrestrial network using ACE was then added to the IT Guru network model for the underway cutter. Some EAs that were not deployed yet or were not accessible on CGDN+ were modeled using generic OPNET templates. Inputs to the model such as number of shipboard users and frequency/duration of EA use are easily modified to conduct "what-if" scenarios based on cutter class and mission. In general, EA usage assumptions were based on how long cutters were away from homeport. Scenarios were conducted for both large (underway more than a week) and small (underway less than a week) cutters.

The Enterprise Applications incorporated into the model included:


- Abstract of Operations (AOPS)
- Coast Guard Human Resource Management System (CGHRMS)
- Coast Guard Message System (CGMS)
- Configuration Management (CMPlus)
- Integrated Aids to Navigation System (I-ATONIS)
- Large Unit Financial System (LUFS NT, LUFS web, and LUFS to go)
- Maritime Information for Safety & Law Enforcement (MISLE)
- Unit Travel System

File Transfer, Email and Web Browsing were also included in the study. Scenarios were run over a 24-hour period for both groups of cutters. In order to baseline existing "pierside" connectivity, initial simulations were run using a T1 (1.5 Mbps) link in the model between ship and shore. During subsequent simulations, the link characteristics were modified to represent a 64 kbps geosynchronous SATCOM (i.e. Inmarsat B) connection for large cutters and a 9600 bps low-earth orbit (LEO) SATCOM (i.e. Iridium or Globalstar) connection for smaller cutters.

Findings

Large cutter:

The following graph shows throughput (bits per second) over a 24-hour period on the large cutter using a T-1 link to shore. Data was generated by using statistical parameters to model expected EA usage onboard ships. Parameters were set for both the applications themselves (i.e. session length) and for user profiles (i.e. number of logons per day and what time). Assumptions used in developing the model can be found in Appendix A.



The table below summarizes the average and peak throughput depicted in the graph above. The maximum throughput in either direction is less than 30 kbps, so it should be safe to assume that the same scenario could be run using a 64 kbps link such as Inmarsat B.

	Average (bps)	Maximum (bps)
Shore to Ship	5592	28360
Ship to Shore	1542	9326

The scenario was run again replacing the T1 link between ship and shore with a 64 kbps geosynchronous SATCOM link. As expected, throughput over the 24-hour period remained about the same because bandwidth was not a limiting factor. The most notable change was an increase in application response time resulting in poor performance and user frustration. Closer analysis indicated that the most important factor for EA performance was latency due to the geosynchronous satellite, approximately 35,600 km above the Earth. The amount of time it takes for a signal to propagate from a transmitter on the ship, up to the satellite and down to an earth station is 238 ms. Thus, a data packet requiring acknowledgement will incur a round trip delay of almost half a second. Many of the Coast Guard applications included in this study required hundreds of round trips, or “application turns” per task. For example, creating a Procurement Request in LUFS NT required 842 application turns. **That adds up to about 400 seconds in**

propagation delay alone! The table below shows the increase in task response time for each application when the satellite link was imposed.

Application	Response Times (sec)		Increase
	T1	Satellite	
AOPS	46.3	65.8	42%
CGHRMS	40.5	76.2	88%
CGMS	13.3	89.3	573%
CMPlus	3.2	36.5	1038%
Email	0.02	0.65	3289%
LUFS	67.7	126.3	87%
MISLE	47.9	60.3	26%
UTS	84.1	127.6	52%
Virus Updates (FTP)	3.2	36.1	1026%
Web Browsing	0.22	3.8	1643%

In order to reduce the effect of satellite latency and improve application performance, we can (1) switch to a low earth orbit (LEO) commercial SATCOM system, (2) create more efficient versions of our EAs for underway cutters and/or (3) optimize network protocols.

(1) Switch to a low earth orbit commercial (LEO) SATCOM system.

- Unlike GEO satellites, which have the same orbital period as the Earth, making them appear “stationary”. LEO systems involve complex networks of satellites that each rotate around the Earth in a few hours. LEO satellites orbit at altitudes between 500 km and 2000 km, reducing propagation delay to 20-25 ms. Although there currently are no LEO systems that can support data rates above 9600 bps to maritime users, recent research done by the Naval Postgraduate School indicates that Qualcomm’s Globalstar system is developing a 128 kbps product for aircraft that could be modified for maritime use if consumer demand warranted it. The authors of this research, LT Kurt Clarke (G-OCC) and LT Andrew Campen (G-SCT), rated Globalstar’s technology in 12 categories along with Inmarsat B and the Navy’s Automated Digital Network Service. Their analysis showed that Globalstar best meets the Coast Guard’s needs. As such, they recommend a partnership be formed with Qualcomm “...to better mold this technology so it can soon become a total solution for the maritime industry.” Another commercial SATCOM system on the horizon is Teledesic, which uses a combination of LEO and medium earth orbit (MEO) satellites offering bandwidths from 128 kbps – 100 Mbps on the uplink and up to 720 Mbps on the downlink. It is touted as the “Internet in the Sky”, but will not be available until 2005.

(2) Create more efficient versions of our EAs for underway cutters.

- The number of application turns per task must be reduced to lessen the effect of satellite latency. Citrix/metaframe applications such as LUFS NT perform the worst since data is sent back and forth with every keystroke or mouse movement. A web-based version of LUFS that is being developed was evaluated in the OPNET model


and was shown to improve performance (response times) even though it used more bandwidth. Another version of the application, called “LUFS to go” is being developed for mobile users and was shown to provide the best performance in the OPNET model. “LUFS to go” puts an Oracle Lite database on every cutter so client and server are both on the shipboard local area network (LAN), resulting in extremely fast response times. Each night, during off-peak hours, the Oracle Lite database is synchronized with the shore-based LUFS database at the USCG Finance Center. LUFS to go appears to be the optimal solution for underway cutters. Other EA program managers should look at this solution as a model to develop cutter-friendly versions of their products.

(3) *Optimize network protocols such as TCP/IP*

- Transmission Control Protocol/Internet Protocol is the “language” computers all over the world use to communicate with each other. These protocols were developed to standardize data transfer on *terrestrial* networks, where congestion and bottlenecks are a problem. On a point-to-point leased commercial satellite link these issues are irrelevant. Inefficient use of the link results when TCP/IP limits packet size, requires acknowledgements of receipt and/or resends duplicate information if the acknowledgement takes too long. Much work has been done by industry to improve TCP/IP performance over satellite links. The R&D Center Advanced Communications Technology Program is sponsoring an ongoing Small Business Innovative Research (SBIR) project to look at CG specific issues in this area.

Small Cutter

The graph below shows throughput (bits per second) over a 24 hour period on a small cutter. Data was generated by using statistical parameters to model expected EA usage onboard ships. Parameters were set for both the applications themselves (i.e. session length) and for user profiles (i.e. number of logons per day and what time). Assumptions used in developing the model can be found in Appendix A.



The table below summarizes the average and peak throughput depicted in the graph above. Note that the maximum throughput in both directions is greater than the 2400 bps, 4800 bps, and 9600 bps presently available commercially.

	Average (bps)	Maximum (bps)
Shore to Ship	3571	19233
Ship to Shore	1000	15118

Smaller cutters will be using Inmarsat Mini-M as the communications path for classified message traffic and will be using HF Messenger to send and receive routine email. Since there aren't many commercially available SATCOM systems to provide CGDN+ connectivity, we need to reassess whether or not cutters underway for less than a week really need access to all of the EAs. Unlike Inmarsat B, for which we are leasing channels on a monthly basis for our large cutters, Inmarsat Mini-M, Globalstar and Iridium channels are not available for lease. Dial-up data connections are inefficient and costly. Some systems are now offering "bandwidth on demand" services, which charge users based on the amount of data sent rather than on length of connection. If it is determined that small cutters require CGDN+ access, a cost study should be conducted to determine the best service plan available.

Conclusion

Based on the findings of this study, although bandwidth is an important issue that must be part of any discussion on cutter connectivity, it is not the only issue of significance. In actuality, it is the latency characteristics that appear to be the limiting factor presently impacting cutter connectivity.

On our large cutters, the present SATCOM link provides up to 64 kbps connectivity. Modeling of present application usage aboard these assets never exceeded 30 kbps, which shows that there is adequate bandwidth for present EAs. However, the design of these applications creates significant delays caused by self-imposed “application turns”. This in turn results in unacceptable performance at the user level due to excessive task response time.

For smaller cutters that presently only have communications links between 2.4 – 9.6 kbps, the problem is a combination of bandwidth limitations and latency delays. Until connectivity capabilities at or above 26.6 kbps are made available to these smaller assets, the Coast Guard must re-evaluate the essential connectivity capabilities we are requiring of the small cutter fleet. Although a reduction of EA’s may result in providing full connectivity to the small cutter fleet, the latency problems encountered by the large cutters will still be present aboard small cutters.

Recommendations

In addition to the recommendations discussed in the findings section of this report (use of alternate (low orbit) communications links, redesign applications to more efficient versions with respect to satellite latency, and optimization of network protocols), the following recommendations should also be investigated:

(1) Optimize EAs for delivery over high-latency satellite link to capital cutters.

Since large cutters are already equipped with Inmarsat B, and the CG is in the process of leasing additional channels to enable full time connectivity, we should take steps to ensure we are making efficient use of this channel. Every developer and/or program manager of USCG Enterprise Applications should be responsible for delivering application content to mobile users with adequate response times. We cannot continue to “jury-rig” our terrestrial applications to make them “work” over a satellite link. We must instead redesign the applications with the end-user in mind. New applications intended to be used by shipboard personnel should be designed to meet performance specifications over the wireless links cutters will have access to.

(2) Look at non-SATCOM alternatives for smaller cutters.

An alternative to full offshore CGDN+ connectivity is to make use of short-range systems such as Cellular Digital Packet Data (CDPD) and 802.11 wireless protocols. CDPD service is provided through cellular telephone companies and offers data rates on the order of 20 kbps, 5-10 miles offshore where cellular towers are present. 802.11 is a line-of-sight system offering much higher speeds of up to 11 Mbps. The R&D Center is currently exploring several applications of this technology, including connectivity in the

A1 coastal zone. If small cutters had CDPD or 802.11 equipment onboard, they would at least be able to come close to shore to download files and access vital applications without having to pull into port and connect.

(3) Use OPNET for application development and follow-on modeling.

The OPNET model was developed as a dynamic tool and is available through the R&D Center for future connectivity studies. New applications could be plugged into the model before they are launched in the field to predict performance and impact on the network. In addition, program managers should be encouraged to use OPNET, or similar products, as a tool in the application design process to ensure the needs of our mobile users are being met.

Appendix A – Profiles

Appendix B - Applications

Appendix A

PROFILES

Profile Name	>wk	<wk	Applications Used	Start Time Offset	Repeat	Duration	Inter-Rep Time
Boarding Officer	3	1	MISLE (BO)	0800-1000	2	End of Last Task	1 – 5 hrs
CGMS_download	1	1	CGMS	0800	4	End of Last Task	Every 4 hrs
CMPlus_Lite_transfer	1	1	CMPlus_lite	0000-0100	1	30 – 60 min	--
Commanding Officer	1	1	AOPS (CO) Email (heavy)	0800-1600 0800-2000	1 -	End of Last Task 1 hr	--
Crewmember	4	2	UTS Email (light) Web-browse (light) Smartforce CGHRMS_user	0800-2300 0000-0100 0000-0100 0800-1000 0800-1000	1 - - 1-3 2	End of Last Task End of Profile End of Profile End of Last Task End of Last Task	-- -- -- 2-4 hrs 2-4 hrs
OPS	1	1	AOPS Email_heavy Web_browsing	0800-1600 0800-2000 0800-2000	1 - -	End of Last Task 1 hr 1 hr	-- -- --
QM	2	1	AOPS	0600-1900	3	End of Last Task	2-4 hrs
SK	3	1	LUFs	0800-1000	3	End of Last Task	2-4 hrs
YN	1	1	CGHRMS	0800-1100	5	End of Last Task	1-3 hrs
Virus Update	1		Virus Update	0006-0017	-	30 – 40 min	--

Appendix B

APPLICATIONS

Application Name	Derivation	Description	Users	Max Throughput	Avg Throughput
AOPS	Trace files collected @ CGC Forward	User logs on. Enters Activity Log. Fills in fields. Logs off.	OPS, QM	716 bps	53.2 bps
AOPS_CO	“”	Same as AOPS, but add “Approve Activity Log” task.	CO		
CGHRMS_admin	“”	User logs on. Checks direct deposit info for several members. Queries for members home address and beneficiary info. Exits.	YN1	27444 bps	1438 bps
CGHRMS_user	“”	User logs on. Creates and submits e-resume. Views e-resume. Modifies own home address. Exits.	crewmember		
CGMS	Trace files collected at R&D Center	Every four hours, ship logs on, downloads all messages and exits. Assume messages are then distributed via ship’s LAN to other users.	CGMS_download	12399 bps	496 bps
CMPlus_Lite	OPNET generic database app	Based on input from G-SLS. Forward store client/server app. Synchronize shipboard dbase w/ server at OSC during low link utilization times.	Engineer	1632 bps	16.3 bps
*(CMPlus/FLS) also tested but not modeled	Data collected at MLCLANT, MLCPAC, OSC	Includes standard FLS tasks than could be integrated into CM_Plus Web app.	N/A	3462 bps	103 bps

Appendix B

Email	OPNET generic apps	Email used throughout the day. Send 3/Get 3 500b msg every hour.	Crewmember		
Email_heavy		Used throughout the day starting after 0800. Send 3/Get 3 2kB msg every 10 min.	CO, OPS	385 bps	109 bps
IATONIS	Simulated using CMPlus upload data	Nightly database synchronization.	QM	1632 bps	16.3 bps
LUFS metaframe	Data collected at RDC & OSC	User logs on, creates 5-10 PRs, approves 2-3 PRs, rejects 1 PR and exits.	SK	7037 bps	607 bps
LUFS web	Data collected at contractor facility on test network	Same sequence as LUFS metaframe, using different input data	SK	15327 bps	1800 bps
LUFS web_lite	“”	Data capture of ftp file containing 1 days worth of ship’s LUFS docs (to simulate dbase synchronization)	SK	1632 bps	16.3 bps

Appendix B

B-2

Appendix B

MISLE_BO	Data collected at RDC using MISLE test network.	User logs on, starts a new activity, fills out all stages of a boarding report, and logs off.	Boarding Officer	16753 bps	1329 bps
MISLE_intel	“”	Used to conduct queries on vessels or personnel. User logs on, enters vessel sighting, searches for person in dbase and logs off.	OPS/CIC watch		
Smartforce	Data collected at RDC using CGDN+	User logs on, personalizes e-learning, launches a course, completes 4 modules and logs off.	crewmember	103632 bps	11382 bps
UTS	Trace files collected @ RDC	User logs on, starts a new travel claim, fills out accounting data, itinerary and expense report. Submits and logs off.	crewmember	2691 bps	108 bps
Virus Update	OPNET generic FTP	Medium-load FTP session, sends 20k file approx every 5 minutes for 30-40min.	Virus software mgr	192 bps	1.94 bps
Web Browsing	OPNET generic HTTP	Medium web browsing throughout the day by 4 crew at any given time. Avg page downloaded every 20 minutes contains 5 small graphics (approx 2-3k)	Crewmember	146 bps	62.6 bps
				** probably too small of an estimate; should be closer to Smartforce #s.	